International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 12-No.1 (2026) pp. 1551-1562
http://www.ijcesen.com

ISSN: 2149-9144

Research Article

Agentic Commerce Applications: How Autonomous AI Is Redefining the Retail & E-Commerce Industry

Sanjay Basu*

TATA Consultancy Services (TCS), USA
* **Corresponding Author Email:** sanjaybasu.work@gmail.com - **ORCID:** 0000-0002-3117-0050

Article Info:

DOI: 10.22399/ijcesen.5091
Received : 02 January 2026
Revised : 01 February 2026
Accepted : 02 February 2026

Keywords

Agentic Artificial Intelligence,
Large Language Models,
E-Commerce Transformation,
Retail Automation,
Autonomous Systems

Abstract:

Retail and e-commerce have steadily evolved from manual operations and rule-based automation to data-driven analytics and, more recently, generative AI systems that assist human decision-making. The next inflection point in this trajectory is agentic artificial intelligence—systems that can interpret goals, plan multi-step actions, coordinate with other agents, and autonomously execute commercial workflows within defined constraints. This paper introduces the concept of agentic commerce, where autonomous AI agents move beyond recommendation and insight generation to actively manage end-to-end retail processes across customer experience, merchandising, operations, and governance.Drawing on recent advances in large language models, multimodal reasoning, and enterprise AI architectures, the paper synthesizes academic literature and industry practices to propose a structured framework for deploying agentic systems in modern retail ecosystems. We examine how agentic capabilities enable continuous demand sensing, dynamic assortment and pricing decisions, conversational and anticipatory shopping experiences, and real-time operational orchestration across omnichannel environments. Beyond technical architecture, the paper addresses organizational, ethical, and governance considerations required to safely operationalize autonomy at scale.By distinguishing assistive AI from truly agentic systems and outlining progressive levels of autonomy in retail decision-making, this work provides both researchers and practitioners with a foundation for understanding how autonomous AI reshapes value creation, competitive advantage, and human roles in digital commerce. The paper concludes by identifying open research challenges and future directions for responsible adoption of agentic commerce in high-frequency retail environments.

1. Scope

The rapid evolution of artificial intelligence in retail has progressed from descriptive analytics and rule-based automation toward predictive modeling, generative content creation, and conversational decision support. While recent advances in large language models (LLMs) and multimodal AI have significantly enhanced retailers' ability to understand customers, optimize assortments, and streamline operations, most deployed systems remain fundamentally assistive in nature—supporting human decision-makers rather than independently executing commercial workflows. As retail ecosystems grow increasingly complex, fast-moving, and omnichannel, the limitations of purely assistive AI become more pronounced.

This paper argues that **agentic artificial intelligence** represents a qualitative shift in the role of AI within retail and digital commerce. Agentic systems move beyond insight generation and recommendation to exhibit characteristics traditionally associated with agency: goal interpretation, multi-step planning, tool invocation, coordination with other agents, autonomous execution, and continuous learning from outcomes. When applied to retail contexts, these capabilities give rise to what this paper defines as **agentic commerce**—a paradigm in which autonomous AI agents increasingly manage end-to-end commercial processes across customer experience, merchandising, supply chain, pricing, and operational orchestration.

The scope of this paper is threefold. **First**, it synthesizes existing academic research and

emerging industry practices to position agentic commerce within the broader historical trajectory of retail technology evolution, distinguishing it clearly from prior waves of automation, analytics, and generative AI. **Second**, it proposes conceptual and architectural frameworks—including a multi-layer LLM-based value chain model and an Agentic Autonomy Framework—that enable systematic reasoning about where, how, and to what extent autonomy should be introduced into retail decision-making. These frameworks are designed to support both scholarly analysis and practical implementation planning. **Third**, the paper grounds theory in practice through detailed use-case analysis, demonstrating how coordinated agents operate across demand sensing, pricing, inventory orchestration, fulfillment, and customer experience in real-world retail scenarios.

Importantly, this work does not frame agentic commerce as a replacement for human expertise. Instead, it examines how increasing levels of autonomy reconfigure human roles—from manual execution toward supervision, policy definition, ethical oversight, and strategic decision-making. The paper explicitly addresses governance, transparency, bias mitigation, accountability, and regulatory compliance as first-class design considerations, recognizing that autonomy without trust undermines both business value and societal acceptance.

From a research perspective, this paper contributes to the literature by bridging gaps between AI systems research, retail operations, and digital commerce strategy. While prior studies have examined recommendation systems, conversational agents, demand forecasting, and supply-chain optimization in isolation, relatively little work has explored how **multi-agent, LLM-driven systems reshape the retail value chain holistically**, particularly under conditions of high decision frequency and real-time adaptation. By articulating progressive levels of agentic autonomy and mapping them to concrete retail use cases, this paper provides a foundation for future empirical research, comparative studies, and regulatory analysis.

The remainder of the paper is structured as follows. Section 2 introduces a theoretical framework for LLM-based value chain transformation and presents an Agentic Autonomy Framework tailored to retail and e-commerce environments. Section 3 provides an end-to-end agentic commerce scenario illustrating coordinated autonomous decision-making in response to a real-world demand shock. Section 4 examines the operational backbone underpinning agentic commerce, spanning inventory management, store operations, fulfillment, and last-mile logistics. Section 5 addresses strategic, ethical, and organizational considerations associated with deploying autonomous AI at scale. Section 6 outlines a future research agenda, highlighting open questions related to agent-to-agent negotiation, regulation, workforce transformation, and consumer trust. The paper concludes by reflecting on the long-term implications of agentic commerce for competitive advantage, human-AI collaboration, and the future structure of digital retail ecosystems.

2. Background and Historical Development of Retail Technology

2.1 Historical Progress in Retail Innovation

The retail industry has witnessed profound technological transformation across four decades, evolving from manually operated point-of-sale systems into deeply instrumented, data-driven digital ecosystems [1]. Early innovations digitized basic transactions, standardized inventory management, and automated store-level reporting. Predictive analytics, recommendation engines, and marketing automation matured over time, yet these systems remained constrained by static rules, narrow task orientation, and limited adaptability to rapidly changing market conditions. Generative AI and multimodal reasoning extended capabilities, enabling natural language interfaces, scalable content production, and conversational shopping experiences [2]. Most enterprise use cases remained fundamentally assistive, where AI supported workers rather than executing work autonomously.

2.2 Understanding Agentic Systems in Commerce

Agentic AI represents a distinct inflection point in retail technology evolution [1]. Traditional AI models react to predefined inputs, whereas agentic systems exhibit characteristics associated with agency, including goal interpretation, multi-step planning, tool invocation, reflection, and autonomous action. These systems execute workflows, coordinate with other agents, and trigger changes in real business systems rather than simply providing recommendations or summarizing data. This shift proves transformational for retail, a sector defined by high-volume decisions, fluctuating demand patterns, and operational complexity spanning supply chain, merchandising, marketing, and customer engagement [2]. Contemporary retailers operate across omnichannel environments where assortment changes occur daily, consumer preferences shift hourly, and

supply chain constraints require real-time adjustments.

2.3 Infrastructure Supporting Autonomous Commerce

Technical enablers of agentic commerce include large language models, multimodal perception systems, vector databases, event-streaming infrastructures, and API-driven commerce systems [3]. These components collectively allow agents to perceive environments, reason about possible actions, and act across digital touchpoints without human intervention. Architectural integration of human-AI interfaces, orchestration layers, simulation environments, and enterprise data platforms within agentic retail ecosystems clarifies how retail systems transition from siloed automation toward interconnected agent networks capable of coordinated decision-making and execution. This architecture demonstrates connections between user interfaces, orchestration layers, simulation sandboxes, and enterprise data platforms that support autonomous commercial behaviors.

2.4 Research Scope and Business Impact

Agentic commerce redefines rather than replaces human expertise [3]. Retail professionals, including merchandisers, planners, marketers, and store associates, shift from manual execution toward supervision, strategy, and oversight of autonomous workflows. Some examples of agentic behavior in retail are virtual shoppers who put together outfits or bundles based on personal preference graphs, autonomous price changes within set limits, and multi-agent systems that work out promotion schedules across channels. Business implications prove significant as retailers adopting agentic systems gain the ability to operate with unprecedented speed, precision, and responsiveness. The shift toward agentic commerce presents challenges, including governance frameworks to ensure safe, aligned behavior; technical integration with legacy systems while maintaining auditability; and organizational role redefinition with workforce reskilling.

3. Theoretical Framework: LLM-Based Value Chain Transformation

3.1 Multi-Layered Architecture Design

Integration of large language models into grocery retail and e-commerce represents systemic transformation of how value is created, delivered, and captured across retail ecosystems [3][4]. The proposed framework positions LLMs as central intelligence layers that orchestrate diverse functions spanning online and physical retail channels. Framework definition comprises four layers: Customer Interaction Layer providing conversational, recommendation-driven, and personalization capabilities; Merchandising and Catalog Layer supporting product categorization, dynamic pricing, and assortment optimization; Operations and Supply Chain Layer enabling autonomous demand sensing, inventory optimization, and delivery decision support; and Governance and Ethics Layer ensuring responsible AI deployment, data privacy, fairness, transparency, and regulatory compliance.

3.2 Language Model Technical Progress

LLMs based on transformer architectures represent substantial leaps in representational flexibility and reasoning capability [3]. Recent frontier models extend capacities by integrating multimodality, including images, audio, video, and structured data, into unified reasoning frameworks. For retail and e-commerce, multimodality proves critical because product understanding relies on diverse signals such as nutritional labels, ingredient lists, packaging images, customer reviews, shelf photos, planograms, and environmental sensor data [4]. Studies show that multimodal LLMs do better than unimodal models at tasks like product categorization, attribute extraction, and recommendation generation. This leads to better customer experiences and operational intelligence.

3.3 Retail-Specific LLM Functionality

Development of agentic LLM systems that plan, decide, and act across tools, APIs, or enterprise workflows represents another major advancement [4]. Such systems autonomously update product metadata, generate substitutions for out-of-stock items, classify large-scale e-commerce catalogs, and perform SKU-level analytics. This autonomy significantly enhances productivity in digital retail environments with vast assortments and frequent content churn. LLMs demonstrate improvements in natural-language search, allowing e-commerce platforms to interpret intent-based queries such as kid-friendly gluten-free snacks or budget meals for families. This capability surpasses keyword-based search used in most traditional e-commerce systems and offers versatility, reasoning depth, and multimodal perception that align with the complexity and data richness of e-commerce ecosystems.

3.4 Knowledge Gaps in Current Literature

Despite rapid progress, several gaps persist in academic literature including limited holistic research on full grocery and e-commerce value chains and insufficient analysis of multimodal LLMs in operational The settings include a scarcity of empirical research on customer-LLM interaction, underdeveloped frameworks for AI governance in high-frequency retail, and a limited examination of LLM agent systems in supply-chain orchestration. A strong academic foundation exists regarding retail AI, deep learning, computer vision, and early conversational systems. However, literature lacks integrative research on how multimodal LLMs reshape combined grocery e-commerce ecosystems across customer, merchandising, and operational domains.

3.5 Agentic Autonomy Framework for Retail and E-Commerce

3.5.1 Rationale for an Autonomy-Based Framework

Much of the existing literature on AI in retail conflates automation, decision support, and autonomy under a single umbrella of "intelligent systems." However, the operational, ethical, and organizational implications of AI differ significantly depending on how much decision authority is delegated to machines. In high-frequency retail environments—where pricing, inventory, promotions, and substitutions change continuously—this distinction becomes critical. To address this gap, we propose an Agentic Autonomy Framework that characterizes retail AI systems based on the degree to which they can independently interpret goals, plan actions, and execute decisions within enterprise constraints.

3.5.2 Levels of Agentic Autonomy in Retail

The proposed framework defines five progressive levels of autonomy, reflecting increasing agentic capability and business impact.

Level 0: Rule-Based Automation
Systems execute predefined rules and workflows without learning or reasoning. Examples include static reorder points, threshold-based promotions, and scripted chatbots. These systems improve efficiency but lack adaptability.

Level 1: Assistive Intelligence
AI models provide insights, predictions, or recommendations to human decision-makers. Examples include demand forecasts, recommendation engines, and LLM-based copilots for merchandising or customer support. Humans retain full control over execution.

Level 2: Tool-Using Agents
AI agents interpret user intent, invoke enterprise tools, and execute bounded actions with human approval or post-action review. Examples include automated catalog enrichment, guided substitutions for out-of-stock items, or AI-assisted markdown recommendations that require managerial validation.

Level 3: Semi-Autonomous Agents
Agents independently plan and execute multi-step workflows within predefined constraints and escalation thresholds. Examples include autonomous inventory rebalancing across stores, dynamic pricing adjustments within guardrails, or promotional orchestration across channels. Human oversight focuses on exception handling rather than routine decisions.

Level 4: Fully Agentic Commerce Systems
Multi-agent systems coordinate across customer experience, merchandising, supply chain, and marketing domains. These agents negotiate trade-offs, learn from outcomes, and adapt strategies over time while aligning with business objectives and governance policies. Humans define strategic goals, ethical boundaries, and risk tolerance rather than operational decisions.

3.5.3 Mapping Retail Use Cases to Autonomy Levels

Not all retail functions require or benefit from full autonomy. High-risk decisions such as regulatory pricing, recalls, or supplier negotiations may remain semi-autonomous, while low-risk, high-volume decisions such as substitutions, content generation, or replenishment can progress more rapidly toward agentic execution. This framework enables retailers to adopt agentic commerce incrementally, aligning autonomy with risk, trust, and organizational readiness.

3.5.4 Human-in-the-Loop and Governance Considerations

Agentic autonomy does not eliminate human involvement; it redefines it. As systems move up the autonomy spectrum, human roles shift from execution toward supervision, policy definition, and exception management. Effective deployment requires clearly defined escalation mechanisms, confidence thresholds, auditability, and explainability to ensure trust and regulatory compliance. The framework therefore integrates governance as a first-class design principle rather than an afterthought.

3.5.5 Research and Design Implications

The Agentic Autonomy Framework provides a foundation for future research into system

evaluation, organizational impact, and ethical design. It enables comparative analysis across retailers and platforms while supporting empirical studies on performance, trust, and long-term competitive dynamics as autonomy increases.

4. Agentic Use Case Deep Dive: An End-to-End Autonomous Retail Scenario

To move beyond conceptual discussion and clearly illustrate how agentic commerce operates in practice, this section presents an end-to-end retail scenario that demonstrates how autonomous agents collaborate, escalate decisions, and learn from outcomes. The example focuses on a common but operationally complex situation in grocery and e-commerce: a sudden spike in demand for a specific product category.

4.1 Scenario Overview: Demand Spike in Frozen Ready Meals

Consider a regional grocery retailer operating both physical stores and an e-commerce platform. During a winter weather event combined with a viral social media trend, demand for frozen ready meals increases sharply within a 48-hour window. Historically, such spikes require manual intervention across forecasting, pricing, inventory allocation, and customer communication—often resulting in stockouts, delayed responses, or inconsistent customer experiences.

In an agentic commerce environment, this situation is handled through coordinated interaction among specialized AI agents, each with clearly defined roles and autonomy boundaries.

4.2 Step 1: Demand Sensing and Situation Awareness

The process begins with a Demand Sensing Agent that continuously monitors multimodal signals across the retail ecosystem. These include historical sales data, real-time e-commerce browsing patterns, search queries, weather forecasts, and external signals such as social media trends. Using large language models integrated with time-series forecasting and event-detection systems, the agent identifies an abnormal surge in intent for frozen ready meals in specific geographic clusters.

Rather than simply generating a forecast report, the agent interprets this signal as a deviation from expected demand patterns and formulates an operational goal: prevent stockouts while maintaining margin and customer satisfaction. This goal is then shared with downstream agents through the orchestration layer.

This capability aligns with recent research on AI-driven demand forecasting in retail supply chains, which highlights the importance of integrating external signals beyond historical sales alone [7].

4.3 Step 2: Pricing and Promotion Coordination

Upon receiving the demand alert, a Pricing Agent evaluates current price elasticity, competitive pricing signals, and inventory depth. Instead of applying static pricing rules, the agent simulates multiple pricing strategies within predefined guardrails set by business leadership. For example, it may recommend holding prices steady for high-demand SKUs to preserve margin while selectively promoting complementary items to increase basket size.

If proposed price adjustments remain within acceptable thresholds, the agent executes changes autonomously across digital channels. However, if margin risk or regulatory sensitivity is detected, the agent escalates the decision to a human pricing manager with a clear explanation of trade-offs and projected outcomes. This human-in-the-loop escalation reflects best practices for responsible AI deployment in high-frequency retail decisions [10].

4.4 Step 3: Inventory Reallocation and Fulfillment Planning

In parallel, an Inventory Orchestration Agent assesses available stock across warehouses, dark stores, and physical retail locations. The agent identifies surplus inventory in lower-demand regions and initiates inter-store transfers while reprioritizing fulfillment routes for online orders. It coordinates with warehouse systems and last-mile delivery platforms to ensure that high-demand areas are replenished first.

For perishable or near-expiry items, the agent collaborates with a Markdown Optimization Agent to apply dynamic discounts or bundle offers, minimizing waste while maintaining service levels. These actions are executed autonomously but logged for auditability and performance review, consistent with emerging governance patterns for AI-driven inventory optimization [8].

4.5 Step 4: Customer Experience and Substitution Management

Despite proactive replenishment, certain SKUs temporarily go out of stock. At this point, a Customer Interaction Agent intervenes across the e-commerce interface and in-store digital touchpoints. Rather than presenting generic out-of-stock messages, the agent generates personalized

substitutions based on customer preferences, dietary constraints, price sensitivity, and past behavior.

For example, a shopper ordering gluten-free frozen meals may receive context-aware alternatives with clear explanations of why each option was recommended. If a customer rejects substitutions repeatedly, the agent adjusts its recommendation strategy in real time. This conversational, adaptive behavior reflects advances in LLM-powered customer experience systems discussed in recent retail AI literature [4][6].

4.6 Step 5: Learning, Reflection, and Continuous Improvement

After the demand spike subsides, a Learning and Evaluation Agent reviews outcomes across the entire workflow. It analyzes metrics such as forecast accuracy, fulfillment rates, substitution acceptance, customer satisfaction, margin impact, and spoilage reduction. Importantly, the agent compares predicted outcomes with actual results and updates internal policies and confidence thresholds accordingly.

Insights from this reflection phase are shared across agents, improving future decision-making. For example, if customers in certain regions consistently reject specific substitutions, the system adjusts its preference models. If pricing actions resulted in unexpected churn, pricing guardrails are recalibrated. This closed-loop learning distinguishes agentic systems from traditional automation and aligns with emerging research on adaptive, self-improving AI architectures [3].

4.7 Why This Scenario Demonstrates True Agentic Commerce

This use case illustrates several defining characteristics of agentic systems in retail. First, agents do not merely respond to isolated inputs; they interpret goals, plan coordinated actions, and execute across multiple systems. Second, autonomy is applied selectively, with built-in escalation for high-risk decisions. Third, learning is continuous and embedded, allowing the system to improve over time rather than repeat static behaviors.

Most importantly, humans remain integral—not as operators executing routine tasks, but as supervisors defining strategy, constraints, and ethical boundaries. This shift in human role reflects the broader organizational transformation required to realize the full value of agentic commerce, as emphasized in both academic and industry research [1][9].

5. Operational Backbone and Distribution Network Enhancement

5.1 Predictive Requirement Analysis

Large language models reshape the operational backbones of grocery retail and e-commerce, enabling intelligent, adaptive, and proactive supply chains and store operations [7]. Beyond customer-facing applications, LLMs function as cognitive decision-support engines that integrate multimodal data streams from point-of-sale transactions and e-commerce browsing behavior to warehouse sensors and in-store cameras to optimize demand forecasting, inventory management, fulfillment, and in-store efficiency. Accurate forecasting proves critical in grocery and e-commerce due to perishable inventory and high SKU diversity, where LLMs analyze historical online orders, browsing data, search queries, and customer preferences to generate high-resolution forecasts at SKU and household levels [7]. Platforms anticipate spikes in demand for frozen vegetables during national holiday weekends or winter, enabling proactive replenishment while integrating POS data, shelf imaging, local weather, and regional events to predict store-level demand.

5.2 Real-Time Stock Management

LLMs support real-time inventory management across e-commerce warehouses, fulfillment centers, and stores [8]. Warehouse optimization involves LLMs coordinating multi-robot picking, packing, and replenishment by integrating inventory, order priority, and delivery schedules, where e-commerce fulfillment centers use AI-enabled systems to route robots and human pickers efficiently. By linking warehouse, store, and online inventory data, LLMs prevent overselling and improve fulfillment accuracy, where shoppers placing online orders can be seamlessly directed to pick up items from the the nearest stores with available stock [8]. For perishable goods management, LLMs recommend shelf rotation, dynamic markdowns, or bundle promotions to reduce spoilage directly impacting profitability and sustainability.

5.3 Brick-and-Mortar Operational Intelligence

LLMs are increasingly applied to physical store operations complementing digital and e-commerce workflows. LLMs analyze real-time shelf images to detect low stock or misplaced items, or planogram deviations where store staff receive actionable instructions for restocking or rearrangement reducing labor-intensive audits. Task prioritization and staff guidance involves LLMs recommending

daily operational priorities such as which aisles require immediate restocking, which promotions need signage updates, and which customer inquiries require assistance. By linking operational insights with customer-facing systems, LLMs enable associates to provide instant guidance on product location, substitutions, and nutritional information. These capabilities demonstrate how LLMs bridge digital intelligence with physical store efficiency, creating seamless omnichannel operational environments.

5.4 Final-Mile Distribution and Cross-Channel Synchronization

E-commerce fulfillment benefits from LLM-driven intelligence in routing, scheduling, and personalization. LLMs optimize delivery routes by integrating order locations, traffic patterns, and time-sensitive constraints, which reduce delivery time and operational costs. Order prioritization and substitutions involve LLMs recommending alternatives for items that are out of stock or suggesting ways to combine multiple orders to make the process more efficient, using AI systems for automatic suggestions and grouped deliveries. The strength of LLMs comes from their ability to connect online, warehouse, and in-store operations by showing a single up-to-date view of inventory across all areas, which helps improve order accuracy and lowers stock errors. Insights from online sales, store transactions, and delivery performance continuously inform forecasting, warehouse routing, and store-level decisions, while LLMs dynamically allocate staff, robotics, and delivery resources based on real-time demand.

6. Business Strategy, Ethics, and Deployment Frameworks

6.1 Information Security and Privacy Compliance

Adoption of large language models in grocery retail and e-commerce introduces profound opportunities while raising critical ethical, governance, and compliance considerations [9]. As LLMs interact with sensitive customer data, operational information, and supplier intelligence, retailers must balance innovation with accountability, transparency, and fairness. LLMs require access to rich data streams, including customer purchase histories, browsing behavior, dietary preferences, and in some cases health-related information such as allergen sensitivities or nutrition tracking [9]. Proper data governance frameworks prove essential to protect privacy and comply with regulations where platforms must obtain explicit consent for data collection and use, allowing users to opt into personalized nutrition recommendations while anonymizing purchase data for broader analytics. Secure data handling requires LLMs to process sensitive data using encryption and secure protocols to prevent unauthorized access, breaches, or misuse while ensuring alignment with GDPR, CCPA, and other relevant privacy regulations.

6.2 Algorithmic Equity and Model Interpretability

LLMs may inadvertently propagate biases present in training data impacting recommendations, promotions, and operational decisions [10]. Models trained on historical purchase data may overrepresent certain brands or demographics, unintentionally disadvantage minority customer segments, while dynamic pricing algorithms must be monitored to prevent discriminatory pricing or unfair targeting. Systems offering higher discounts only to certain zip codes may inadvertently reinforce inequities requiring continuous evaluation, bias audits, and model retraining with diverse datasets to ensure equitable outcomes [10]. LLMs operate as complex black-box models, making explainability crucial for operational trust and customer confidence. Shoppers interacting with recommendations or substitutions should understand why certain products are suggested while store managers, warehouse operators, and e-commerce staff must interpret LLM-generated instructions for stocking, routing, and promotions, where transparent reasoning supports faster human validation and error correction.

6.3 Market Advantage and Positioning Strategies

Emergence of large language models represents structural shifts in how grocery retailers and e-commerce platforms compete. Unlike previous waves of digital transformation such as mobile apps, loyalty programs, or cloud analytics, LLMs influence not just channels but core value propositions, enabling new forms of personalization, automation, and operational intelligence. Retailers increasingly differentiate not by assortment breadth alone but by fit between customer needs and product offerings, where LLMs allow shifts from segment-level targeting to individual-level understanding synthesizing dietary habits, cultural preferences, health goals, and real-time intent. Grocers create hyper-personalized meal planning based on budget, time constraints, and health considerations, while e-commerce platforms

offer context-aware product bundles and retailers with more advanced recommendation engines observe measurable gains in basket size, conversion rates, and repeat purchase probability.

6.4 Deployment Roadmap and Organizational Change

Successful deployment of large language models in grocery retail and e-commerce requires more than technical readiness; it demands organizational alignment, workforce enablement, cross-functional governance, and scalable operating models. Empirical research in digital transformation emphasizes that AI initiatives fail not due to algorithmic shortcomings but due to inadequate integration with business processes, cultural resistance, and limited executive sponsorship. LLM transformation begins with clarity on strategic intent, where retailers leading in AI adoption typically articulate enterprise-level AI charters that define organizational innovation priorities, ethical principles, data responsibilities, and acceptable risk thresholds. This vision should specify customer-facing goals such as hyper-personalization and frictionless commerce, operational goals such as intelligent forecasting and automated task orchestration, and e-commerce goals including conversational search and dynamic catalog management.

6.5 Workforce Adaptation and Structural Redesign

Retailers must operationalize vision through federated AI operating models in which central AI teams manage governance, architecture, privacy, and model evaluation while business units co-develop domain-specific LLM use cases and technology teams ensure platform scalability and integration with enterprise systems. This structure enables agility while preserving model safety and reliability. LLMs unlock value only when built on robust, well-governed data ecosystems, where grocery retailers often grapple with fragmented product catalogs, inconsistent SKU hierarchies, and siloed customer data conditions that undermine model accuracy. Scalable LLM foundations require unified product knowledge graphs that harmonize product metadata, nutrition profiles, brand hierarchies, dietary tags, and supply-chain attributes, enabling LLMs to reason about product substitution, recipe constraints, and dietary suitability. Additionally, unified customer data layers integrating transactional, behavioral, loyalty, and digital signals ensure privacy-compliant foundations for personalized recommendations and conversational engagement.

7. Future Research Agenda for Agentic Commerce

While agentic commerce presents significant opportunities to reshape retail and e-commerce, it also introduces complex technical, regulatory, and societal questions that remain underexplored in current literature. As autonomous AI systems increasingly act on behalf of consumers, retailers, and brands, future research must move beyond isolated use cases toward a deeper understanding of systemic, long-term implications. This section outlines four priority research directions essential for advancing agentic commerce in a responsible, scalable, and trustworthy manner.

7.1 Agent-to-Agent Negotiation and Market Dynamics

One of the most transformative—and least studied—aspects of agentic commerce is the emergence of agent-to-agent interaction. In future retail ecosystems, autonomous agents representing consumers, retailers, suppliers, and media platforms may negotiate prices, delivery windows, trade promotions, and advertising spend in real time. Unlike traditional algorithmic pricing systems, these agents operate with strategic intent, adaptive learning, and negotiation capabilities powered by large language models and reinforcement learning.
This raises fundamental research questions. How should negotiation protocols be designed to prevent manipulation, collusion, or unintended market instability? What happens when consumer agents optimize aggressively for price while retailer agents optimize for margin and inventory health? Game-theoretic modeling combined with agent-based simulations offers promising directions for understanding emergent behaviors in such environments [11]. Without rigorous study, agent-driven marketplaces risk amplifying volatility rather than improving efficiency.

7.2 Regulation and Governance of Autonomous Commerce

As autonomy increases, existing regulatory frameworks struggle to keep pace. Current retail regulations assume human decision-makers, clear accountability, and deterministic systems. Agentic commerce challenges these assumptions by introducing systems that act continuously, adapt over time, and coordinate across organizational boundaries.

Future research is needed to define regulatory models for autonomous commercial decision-making, including accountability structures, audit requirements, and certification standards for agentic systems [12]. Key questions include how liability should be assigned when autonomous agents make harmful decisions, how regulators can audit reasoning processes within opaque models, and how compliance can be enforced across dynamic, learning systems. Cross-disciplinary research involving law, AI ethics, and retail operations is essential to develop governance frameworks that protect consumers while allowing innovation to flourish.

7.3 Long-Term Workforce Impact and Organizational Transformation

While much discussion focuses on productivity gains, the long-term impact of agentic commerce on the retail workforce remains insufficiently studied. As autonomous agents take over routine decision-making and execution tasks, human roles shift toward supervision, strategy, exception handling, and ethical oversight. This transition fundamentally alters skill requirements, career paths, and organizational structures.

Future research should examine how different levels of agentic autonomy affect workforce satisfaction, decision quality, and organizational resilience over time. Longitudinal studies could explore how retailers reskill employees, redefine accountability, and balance human judgment with machine autonomy. Understanding these dynamics is critical, as poorly managed transitions risk workforce disengagement and loss of domain expertise, undermining the very benefits agentic systems promise to deliver.

7.4 Consumer Trust in AI Intermediaries

As AI agents increasingly act as intermediaries between consumers and retailers—selecting products, making substitutions, and optimizing purchases—consumer trust becomes a central determinant of adoption and long-term success. Unlike traditional recommendation systems, agentic systems do not merely suggest options; they make decisions on behalf of users, often with limited visibility into underlying reasoning.

Future research must explore how transparency, explainability, and user control influence trust in agentic commerce [13]. Questions include how much autonomy consumers are willing to delegate, how trust evolves over repeated interactions, and how trust breaks down when agents make mistakes. Experimental studies examining different explanation styles, consent mechanisms, and override options can inform the design of AI systems that enhance trust rather than erode it. Without trust, even the most technically advanced agentic systems risk rejection or regulatory backlash.

7.5 Toward a Responsible and Sustainable Agentic Commerce Ecosystem

Collectively, these research directions highlight that agentic commerce is not merely a technological shift but a systemic transformation of retail ecosystems. Progress requires coordinated advances in AI system design, economic modeling, regulatory frameworks, organizational theory, and human–computer interaction. By addressing these open questions, future scholarship can help ensure that agentic commerce delivers sustainable value while preserving fairness, accountability, and human agency in increasingly autonomous retail environments.

Table 1: *Evolution of Retail Technology Phases [1][2]*

Phase	Era	Key Technologies	Capabilities	Limitations
Manual Processes	1980s-1990s	Point-of-sale terminals, Barcode scanners, Early ERP systems	Basic transaction digitization, Inventory tracking, Centralized records	Human-driven decisions, Periodic reporting, Reactive operations
Analytical Automation	2000s-2010s	Data warehouses, Business intelligence, Predictive analytics	Demand forecasting, Customer segmentation, Automated reporting	Static rules, Limited adaptability, Narrow task focus
Generative AI	2015-2023	Machine learning, Recommendation engines, NLP interfaces	Personalized recommendations, Content generation, Conversational interfaces	Assistive only, Human execution required, Limited autonomy

Agentic Systems	2024-Present	Large language models, Multimodal AI, Autonomous agents	End-to-end workflow execution, Multi-agent coordination, Autonomous decision-making	Governance challenges, Integration complexity, Skill gaps

Table 2: *Four-Layer LLM Framework for Retail Transformation [3] [4]*

Layer	Primary Functions	Key Capabilities	Business Impact
Customer Interaction	Conversational commerce, Personalized recommendations, Digital navigation	Natural language queries, Intent understanding, Context-aware suggestions, Multi-turn dialogues	Increased engagement, Higher conversion rates, Enhanced satisfaction, Larger basket sizes
Merchandising & Catalog	Product classification, Dynamic pricing, Assortment optimization, Content generation	Automated SKU categorization, Attribute extraction, Price elasticity analysis, Promotional planning	Faster catalog updates, Improved accuracy, Reduced manual effort, Optimized margins
Operations & Supply Chain	Demand forecasting, Inventory management, Warehouse coordination, Delivery optimization	Predictive analytics, Real-time stock visibility, Robot coordination, Route planning	Reduced stock-outs, Lower spoilage, Faster fulfillment, Cost efficiency
Governance & Ethics	Privacy protection, Bias mitigation, Transparency, Regulatory compliance	Data encryption, Fairness audits, Explainable AI, Policy enforcement	Trust building, Risk mitigation, Legal compliance, Brand protection

Table 3: *LLM Applications in Customer Experience Enhancement [4][5]*

Application Area	LLM Functionality	Customer Benefits	Retail Outcomes
Hyper-Personalization	Synthesizes purchase history, dietary preferences, budget constraints, browsing patterns	Tailored product suggestions, Balanced meal plans, Dietary compliance support	Increased basket size, Higher customer lifetime value, Improved retention
Conversational Search	Interprets complex natural language queries, Maintains context across interactions	Quick product discovery, Intent-based results, Refined search capabilities	Reduced search abandonment, Faster purchase decisions, Enhanced satisfaction
Smart Substitutions	Analyzes product attributes, nutritional profiles, price points, availability	Equivalent alternatives, Dietary-appropriate options, Budget-conscious choices	Maintained order fulfillment, Reduced cancellations, Customer trust
Recipe Integration	Links ingredients with recipes, Suggests complementary items, Provides cooking guidance	Meal planning assistance, Nutritional balance, Convenient shopping	Cross-category purchases, Increased units per transaction, Repeat visits
Proactive Nudges	Monitors contextual signals, Analyzes seasonal patterns, Tracks promotional cycles	Timely reminders, Relevant promotions, Fresh produce alerts	Higher promotional uptake, Seasonal sales optimization, Reduced waste

Table 4: *Ethical and Governance Considerations for LLM Deployment [9] [10]*

Consideration	Key Challenges	Mitigation Strategies	Regulatory Alignment

Area			
Data Privacy	Sensitive customer information, Health data exposure, Purchase history tracking, Behavioral profiling	Explicit consent mechanisms, Data anonymization, Encryption protocols, Access controls, Data minimization	GDPR compliance, CCPA requirements, Health data regulations, Cross-border data transfer laws
Algorithmic Bias	Historical data biases, Demographic underrepresentation, Price discrimination, Unfair recommendations	Diverse training datasets, Continuous bias audits, Fairness metrics, Regular model retraining, Stakeholder feedback	Anti-discrimination laws, Fair pricing regulations, Consumer protection statutes
Transparency	Black-box models, Unexplainable recommendations, Opaque pricing logic, Hidden decision factors	Explainable AI techniques, Customer-facing explanations, Audit trails, Decision documentation, Transparent policies	Consumer right-to-explanation, Algorithmic accountability laws, Truth in advertising
Security	Data breaches, Unauthorized access, Model manipulation, Adversarial attacks	Secure infrastructure, Regular security audits, Penetration testing, Incident response plans, Access monitoring	Data breach notification laws, Cybersecurity standards, Industry compliance frameworks
Accountability	Autonomous decision errors, Liability assignment, System failures, Customer harm	Human oversight mechanisms, Escalation protocols, Clear ownership structures, Insurance coverage, Redress processes	Product liability laws, Consumer protection regulations, Professional standards

8. Conclusions

Agentic artificial intelligence represents transformational paradigms in retail, extending beyond assistive technologies to autonomous systems capable of end-to-end execution across commercial workflows. Integration of large language models across customer experience, merchandising, and operations demonstrates how multimodal reasoning, natural language understanding, and autonomous planning capabilities enable retailers to operate with unprecedented speed, precision, and personalization. Four-layer frameworks encompassing customer interaction, merchandising, operations, and governance provide comprehensive architectures for understanding how LLMs reshape value creation and delivery in modern retail ecosystems. Building durable competitive advantage through LLM adoption requires balancing innovation with responsible governance while establishing robust data foundations and organizational alignment. Retailers must articulate clear enterprise AI visions, implement federated operating models that preserve agility while ensuring safety, and invest in unified data ecosystems that enable accurate multimodal reasoning. Phased adoption roadmaps from foundational capabilities through differentiated customer value to ecosystem leadership provide practical frameworks for organizations at various stages of digital maturity. Organizations that reach advanced integration levels build structural advantages through proprietary models, seamless customer experiences, and comprehensive AI integration across media networks, logistics, and merchandising. Several critical areas warrant continued scholarly attention, including empirical studies on customer-LLM interactions in real-world grocery and e-commerce settings, longitudinal analysis of competitive dynamics as agentic systems become more prevalent, and development of comprehensive regulatory frameworks for ethical AI deployment in high-frequency retail environments. Additional research is needed on organizational change management requirements for successful agentic commerce adoption, evolving role definitions for retail professionals in AI-augmented environments, and technical architectures that best support human-in-the-loop escalation mechanisms while maintaining autonomous execution efficiency. Agentic

commerce represents both an inevitability and a transformative opportunity, reshaping the structure and operation of digital retail ecosystems. As autonomous AI systems increasingly become intermediaries acting on behalf of shoppers, merchandisers, and marketers, the retail industry will experience fundamental shifts in how value is created, competitive advantage is sustained, and human expertise is applied. Successful integration of agentic AI requires not only technical sophistication but also ethical commitment, organizational readiness, and strategic vision to ensure that autonomous systems enhance rather than diminish human elements that remain central to retail excellence.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript.

References

[1] Shervin Ghaffari, et al., "Generative-AI in E-Commerce: Use-Cases and Implementations," IEEE Access, 25 March 2024. Available: https://ieeexplore.ieee.org/document/10475266?utm_source=copilot.com

[2] Ananya Ghosh Chowdhury, et al., "Revolutionizing Retail Operations through Generative AI: A Systematic Review," International Journal of Computer Applications (IJCA), May 2025. Available: https://www.ijcaonline.org/archives/volume187/number4/chowdhury-2025-ijca-924834.pdf?utm_source=copilot.com

[3] Mohaimenul Azam Khan Raiaan, et al., "A Review on Large Language Models: Architectures, Applications, Taxonomies, Open Issues and Challenges," IEEE Transactions on Neural Networks and Learning Systems, 13 February 2024. Available: https://ieeexplore.ieee.org/document/10433480?utm_source=copilot.com

[4] Farooq Shareef, et al., "RetailGPT: A Fine-Tuned LLM Architecture for Customer Experience and Sales Optimization," IEEE Access, 28 November 2024. Available: https://ieeexplore.ieee.org/abstract/document/10760685?utm_source=copilot.com

[5] Aishwarya Gowda A G, et al., "Personalized E-commerce: Enhancing Customer Experience through Machine Learning-driven Personalization," IEEE ICITEICS Proceedings, 22 August 2024. Available: https://ieeexplore.ieee.org/document/10624901?utm_source=copilot.com

[6] Reshmi Tatikonda, et al., "Chatbot and its Impact on the Retail Industry," IEEE Transactions on Services Computing, 13 March 2025. Available: https://ieeexplore.ieee.org/abstract/document/10915098?utm_source=copilot.com

[7] J. Li, et al., "Automated Demand Forecasting in Retail Supply Chains Using AI," IEEE Transactions on Engineering Management, May 2025. Available: https://ieeexplore.ieee.org/abstract/document/11154572?utm_source=copilot.com

[8] Shreyas Bailkar, et al., "Smart Inventory Optimization using Machine Learning Algorithms," IEEE Access, 22 March 2024. Available: https://ieeexplore.ieee.org/document/10467512?utm_source=copilot.com

[9] Rakesh Reddy Charla, et al., "Federated Data Engineering for Privacy-Aware AI: Patterns from Distributed Retail and Financial Workflows," IEEE Transactions on Big Data, September 2025. Available: https://ieeexplore.ieee.org/document/11118887?utm_source=copilot.com

[10] Dongsoo Moon, et al., "Metrics and Algorithms for Identifying and Mitigating Bias in AI Design: A Counterfactual Fairness Approach," IEEE Transactions on Artificial Intelligence, 31 March 2025. Available: https://ieeexplore.ieee.org/document/10945860?utm_source=copilot.com

[11] David C. Parkes and Michael P. Wellman, "Economic Reasoning and Artificial Intelligence," *Science*, 17 July 2015. Available: https://www.science.org/doi/10.1126/science.aac8103

[12] European Commission, "Artificial Intelligence Act: Regulatory Framework for AI Systems," 2024. Available: https://artificialintelligenceact.eu/

[13] Ben Shneiderman, "Human-Centered AI: Reliable, Safe & Trustworthy Systems," *ACM Interactions*, 2020. Available: https://dl.acm.org/doi/10.1145/3375627